

Claims

1. A method of routing traffic from a source node to a destination node in a mesh topology network connected to a plurality of hosts, said network having a plurality of nodes that are routers, including source node and destination node and a plurality of links connecting said nodes, the method comprising the steps of:

computing a node metric for each node of said plurality of nodes, said node metric reflecting the congestion level of said node by accounting for future scheduled traffic from-neighboring nodes and to-neighboring nodes;

determining a path metric for each path of a plurality of paths from source node to destination node; and

allocating the load from source node to destination node to said plurality of paths according to the path metric of each path of said plurality of paths.

2. The method of claim 1 wherein said computing a node metric for each node comprising the steps of:

calculating a from-neighbor component of said node metric, said from-neighbor component reflecting the future traffic load from a plurality of neighbors of said node to said node;

calculating a to-neighbor component of said node metric, said to-neighbor component reflecting the future traffic load from said node to said plurality of neighbors; and

combining said from-neighbor component with said to-neighbor component to yield said node metric.

3. The method of claim 2 wherein said calculating a from-neighbor component of said node metric comprises the steps of:

determining a first future traffic load for each link of a first plurality of links of said node; and
combining the first future traffic load of said first plurality of links of said node to yield said from-neighbor component of said node metric.

4. The method of claim 3 wherein said determining a first future traffic load for each link of a first plurality of links of said node comprises the steps of:

determining the length of a first queue in a first neighbor at the other end of said link, said first queue storing the packets to be sent to said node;

determining the available bandwidth of said first neighbor for sending data to said node; and

analyzing said length of said first queue and said available bandwidth of said first neighbor to yield said first future traffic load for said link.

5. The method of claim 4 wherein said analyzing also takes into account a bandwidth granted by said node to said first neighbor.

6. The method of claim 4 wherein said analyzing comprises using an interpolation method.

7. The method of claim 3 wherein said calculating a to-neighbor component of said node metric comprises the steps of:

determining a second future traffic load for each link of a second plurality of links of said node;
and

combining the second future traffic load of said second plurality of links of said node to yield said to-neighbor component of said node metric.

8. The method of claim 2 wherein said calculating a to-neighbor component of said node metric comprises the steps of:

determining a second future traffic load for each link of a second plurality of links of said node;
and

combining the second future traffic load of said second plurality of links of said node to yield said to-neighbor component of said node metric.

9. The method of claim 8 wherein said determining a second future traffic load for each link of a second plurality of links of said node comprises the steps of:

determining the length of a second queue in said node, said second queue storing the packets to be sent to a second neighbor being at the other end of said link;

determining the available bandwidth of said node for sending data to said second neighbor; and

analyzing the length of said second queue and said available bandwidth of said node to yield said second future traffic load for said link.

10. The method of claim 9 wherein said analyzing also takes into account a bandwidth granted by said node to said second neighbor.

11. The method of claim 9 wherein said analyzing comprises using an interpolation method.

12. A method of routing traffic load from a source node to a destination node in a mesh topology network connected to a plurality of hosts, said network having a plurality of nodes that are routers, including source node and destination node and a plurality of links connecting said nodes, the method comprising the steps of:

computing a node metric for each node of said plurality of nodes by calculating a metric contribution for each link of a first plurality of links of said node, accounting for future scheduled traffic to and from the node in each link, and combining said metric contributions of said first plurality of links to yield said node metric;

determining a path metric for each path of a plurality of paths from source node to destination node; and

allocating the load from source node to destination node to said plurality of paths according to the path metric of each path of said plurality of paths.

13. The method of claim 12 wherein said calculating a metric contribution for each link comprises the steps of:

determining a first future traffic load for said link from a neighbor being at the other end of said link to said node;

determining a second future traffic load for said link from said node to said neighbor; and

combining said first future traffic load and said second future traffic load to yield said metric contribution for said link.

14. The method of claim 13 wherein said determining a first future traffic load comprises the steps of:

determining the length of a first queue in said neighbor, said first queue storing the packets to be sent to said node;

determining the available bandwidth of said neighbor for sending data to said node; and

analyzing said length of said first queue and said available bandwidth of said neighbor to yield said first future traffic load for said link.

15. The method of claim 14 wherein said analyzing also takes into account a bandwidth granted by said node to said neighbor.

16. The method of claim 14 wherein said analyzing comprises using an interpolation method.

17. The method of claim 14 wherein said determining said second future traffic load comprises the steps of:

determining the length of a second queue in said node, said second queue storing the packets to be sent to a said neighbor;

determining the available bandwidth of said node for sending data to said neighbor; and

analyzing the length of said second queue and said available bandwidth of said node to yield said second future traffic load for said link.

18. The method of claim 13 wherein said determining said second future traffic load comprises the steps of:

determining the length of a second queue in said node, said second queue storing the packets to be sent to a said neighbor;

determining the available bandwidth of said node for sending data to said neighbor; and

analyzing the length of said second queue and said available bandwidth of said node to yield said second future traffic load for said link.

19. The method of claim 18 wherein said analyzing also takes into account a bandwidth granted by said neighbor to said node.

20. The method of claim 18 wherein said analyzing comprises using an interpolation method.